

RIDS: Robust Identification of Sparse Gene Regulatory Networks from Perturbation Experiments

Hoi-To Wai, Anna Scaglione, Uzi Harush, Baruch Barzel and Amir Leshem*

Abstract

Reconstructing the causal network in a complex dynamical system plays a crucial role in many applications, from sub-cellular biology to economic systems. Here we focus on inferring gene regulation networks (GRNs) from perturbation or gene deletion experiments. Despite their scientific merit, such perturbation experiments are not often used for such inference due to their costly experimental procedure, requiring significant resources to complete the measurement of every single experiment. To overcome this challenge, we develop the Robust IDentification of Sparse networks (RIDS) method that reconstructs the GRN from a small number of perturbation experiments. Our method uses the gene expression data observed in each experiment and translates that into a steady state condition of the system’s nonlinear interaction dynamics. Applying a sparse optimization criterion, we are able to extract the parameters of the underlying weighted network, even from very few experiments. In fact, we demonstrate *analytically* that, under certain conditions, the GRN can be perfectly reconstructed using $K = \Omega(d_{max})$ perturbation experiments, where d_{max} is the maximum in-degree of the GRN, a small value for realistic sparse networks, indicating that RIDS can achieve high performance with a scalable number of experiments. We test our method on both synthetic and experimental data extracted from the DREAM5 network inference challenge. We show that the RIDS achieves superior performance compared to the state-of-the-art methods, while requiring as few as $\sim 60\%$ less experimental data. Moreover, as opposed to almost all competing methods, RIDS allows us to infer the directionality of the GRN links, allowing us to infer empirical GRNs, without relying on the commonly provided list of transcription factors.

1 Introduction

Our ability to understand, predict and manipulate the behavior of sub-cellular systems is limited by our currently partial maps of the cell’s biochemical process, obstructing the exposure of microscopic biological mechanisms and hindering drug development [1,2,7,10,18]. To make advances, researchers rely on high-throughput profile technologies, such as DNA microarrays, where a large amount of gene expression data are used for reverse engineering gene regulatory networks (GRNs). The scale of these experiments allows one to infer the GRN by extracting the correlations [21] or other related statistical similarity measures [11, 20] in the expression levels genes under the different experimental conditions. Alternative methods rely on techniques such as feature selection, time series regression, etc., [6, 13, 16, 17, 22, 23, 27, 28, 31, 32]. The problem is that such methods do not directly infer the causal relationships between genes, often falsely identifying indirect correlations

*H.-T. Wai and A. Scaglione are with School of ECEE, Arizona State University, USA. Emails: {htwai, Anna.Scaglione}@asu.edu. U. Harush and B. Barzel are with Department of Mathematics, Bar-Ilan University, Israel. Emails: {uziharush, baruchbarzel}@gmail.com. A. Leshem is with Faculty of Engineering, Bar-Ilan Univer-

as physical interactions [3]. Another challenge is that statistical similarity measures are symmetric, and hence lack information on the directionality of the links. Therefore, one must rely on exogenous information, such as a predefined list of transcription factors [15,21], without which it is not possible to deduce the direction of the inferred links.

To address these challenges, researchers attempt to infer the structure of the GRN directly from genetic perturbation data [6,8,12,19,22,26,32]. In each experiment one or more genes is perturbed, *e.g.*, knocked out or over/under-expressed, and the change in the expression of all remaining genes is observed. Using perturbation data has two main advantages: (i) as opposed to correlations and other related measures, that may result from spurious statistical dependencies, perturbations represent a controlled experimental measure of causal effects [5, 19]; (ii) perturbations provide directional information, as i will only respond to j if there is a directed link/path leading from j to i . The problem is that perturbation experiments are of limited scale, with the most comprehensive datasets incorporating at most several hundreds of perturbations, orders of magnitude less than the $\sim 10^4$ genes in a typical GRN. Hence, despite their crucial advantages, we face severe limitations in our ability to use perturbation data for sub-cellular network inference.

Here we exploit the fact that real GRNs are extremely sparse, with the typical gene regulated by at most a few other genes. To take advantage of this, we develop the Robust IDentification of Sparse networks (RIDS) method, that allows us to infer GRN from a small number of perturbations. Our method takes as input a set of perturbation experiments, and as output provides a minimal (sparse) network $\{A_{ij}\}_{i,j=1}^n$ that is consistent with the observed data. Testing our approach against the DREAM5 gold standard [21] we find that with as few as $\sim 35\%$ of the total input provided in DREAM5, we are able to match, and in certain cases, even surpass the top performing inference methods. Since perturbations also encode information on the direction of the causal influence relationship, RIDS allows us to infer the directionality of the links, as we demonstrate by applying our method without using the DREAM5 provided list of transcription factors. Strikingly, we find that even without the transcription factor list, and despite using significantly less input data, RIDS remains competitive with the top performing methods of DREAM5. Finally, an additional advantage that our perturbation-based inference provides is that the inferred link weights A_{ij} capture the true strength of the interactions, providing an estimate for the rate constants for the different genetic reactions. This is in contrast to the majority of inference methods that only provide likelihood estimates, that, on their own, bare no inherent meaning, other than indicating the likelihood for the existence or in-existence of a link.

Notations. For any natural number $n \in \mathbb{N}$, we denote $[n]$ as the set $\{1, 2, \dots, n\}$. Vectors (*resp.* matrices) are denoted by boldfaced letters (*resp.* capital letters). We denote x_i as the i th element of the vector \mathbf{x} , $[\mathbf{E}]_{\mathcal{S},:}$ (*resp.* $[\mathbf{E}]_{:, \mathcal{S}}$) denotes the submatrix of $\mathbf{E} \in \mathbb{R}^{m \times n}$ with only the rows (*resp.* columns) selected from $\mathcal{S} \subseteq [m]$ (*resp.* $\mathcal{S} \subseteq [n]$). Vector $\mathbf{e}_k \in \mathbb{R}^n$ is a unit vector with zeros everywhere except for the k th coordinate. The superscript $(\cdot)^\top$ denotes matrix/vector transpose. $\|\cdot\|_2$ denotes the Euclidean norm and $\|\cdot\|_1$ is the ℓ_1 -norm.

2 Method & Main Results

The GRN of interest is modeled as a directed graph $G = (V, E)$, where $V = [n] = \{1, \dots, n\}$ is the set of genes and $E \subseteq [n] \times [n]$ is the edge set. Gene j is said to be a regulator (or deregulator) of gene

sity, Israel. Email: leshem.amir2@gmail.com.

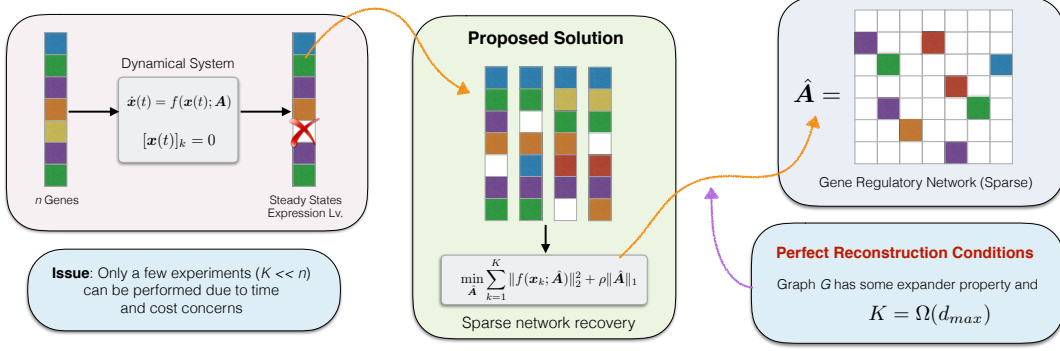


Figure 1: Conceptual overview of RIDS. The proposed method takes the *steady state* gene expression levels resulted from K ($K \ll n$ where n is the number of genes) distinct perturbation experiments and tackle a sparse optimization problem to recover the GRN. For perfect recovery, under some conditions on the graph structure, the proposed solution requires only $K = \Omega(d_{max})$ experiments, where d_{max} is the maximum in-degree of the GRN.

i if $(i, j) \in E$, i.e., there exists an edge from j to i . The GRN is associated with a weight matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ that encodes the strengths of regulation between the genes. We have $[\mathbf{A}]_{ij} = A_{ij} = 0$ if and only if $(i, j) \notin E$. Here we assume the absence of self-links, hence $\text{diag}(\mathbf{A}) = \mathbf{0}$. The time dependent expression level of gene i , $x_i(t)$, follows the nonlinear dynamic:

$$\dot{x}_i(t) = -x_i(t) + \sum_{j=1}^n A_{ij} \cdot h(x_j(t); \mathbf{b}), \quad \forall i \in [n]. \quad (1)$$

where $\dot{x}_i(t) := dx_i(t)/dt$ is the rate of change of the expression level $x_i(t)$. The first term on the right hand side captures the gene i 's self dynamics, capturing processes such as degradation, and the sum captures the impact of i 's interacting partners; $h(x; \mathbf{b})$ is the nonlinear continuous response function describing the regulatory mechanism such that \mathbf{b} describes its parameters. For instance, setting $h(x) = c \cdot x^b$ describes chemical activation, where according to the law of mass action b is the level of cooperation in the activating process. Another frequently used response function is the Hill function $h(x) = x^b / (1 + x^b)$, a saturating function such that $\lim_{x \rightarrow \infty} h(x) = 1$, which captures a *switch-like* activation process.

In matrix form, we can express any *unperturbed* steady state $\bar{\mathbf{x}}$ of the system (1) as:

$$\bar{\mathbf{x}} = \mathbf{A} \mathbf{h}(\bar{\mathbf{x}}; \mathbf{b}), \quad (2)$$

where $\mathbf{h}(\mathbf{x}; \mathbf{b}) := (h(x_1; \mathbf{b}), h(x_2; \mathbf{b}), \dots, h(x_n; \mathbf{b}))^\top$ is a column vector. Next, we consider a set of K distinct perturbation experiments. In each experiment $k = 1, 2, \dots, K$, we *fix* the state of gene k at a desired value $z_k \in \mathbb{R}$, i.e., $x_k(t) = z_k$ for all t . The perturbed steady state $\mathbf{x}[k]$ takes the form

$$\mathbf{x}[k] = (\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \mathbf{h}(\mathbf{x}[k]; \mathbf{b}) + z_k \mathbf{e}_k. \quad (3)$$

Notice that when $z_k = 0$, this corresponds to the deletion of gene k in the experiment.

Suppose that the parameter \mathbf{b} is known. In Eq. (3), the values of all $\mathbf{x}[k]$ are extracted from the experimental results, hence despite the nonlinear interactions, the equation is linear in the unknown \mathbf{A} . As for each k , Eq. (3) constitutes a set of $(n - 1)$ -linear equations in the unknown parameter \mathbf{A} . Naively, we must conduct perturbation experiments for all n genes, yielding the required $n(n - 1)$ equations to reconstruct the $n \times (n - 1)$ off-diagonal terms of the unknown GRN. Such comprehensive perturbation experiments, however, are seldom available. Indeed, the GRN

of most organisms comprises $n \sim 10^3$ genes, far exceeding the scale of the majority of microarray experiments, which, given the level of available resources, consists of $K \sim 10^1 - 10^2$ experiments. Hence, we shall focus on the limit where $K \ll n$ and derive a sparse optimization algorithm to extract \mathbf{A} from the resulting underdetermined linear system (3).

We observe a curious property that can help revealing the structure of \mathbf{A} . Denote $\mathbf{a}_{col,k}$ as the k th column vector of \mathbf{A} . Let $\nabla \mathbf{h}(\mathbf{x}[k]; \mathbf{b})$ be the diagonal matrix with the i th diagonal element being $h'(x_i[k]; \mathbf{b})$, i.e., the derivative of $h(x_i[k]; \mathbf{b})$ with respect to $x_i[k]$ evaluated at $x_i[k]$, we have:

Proposition 1 *Consider the dynamics (1). Assume that the perturbation in the steady states for the k th experiment, $\bar{\mathbf{x}} - \mathbf{x}[k]$, is small and $\lambda_{max}((\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \nabla \mathbf{h}(\mathbf{x}[k]; \mathbf{b})) < 1$. The perturbation in the steady states can be approximated by:*

$$\bar{\mathbf{x}} - \mathbf{x}[k] \approx ([\bar{\mathbf{x}}]_k - z_k) \mathbf{e}_k + ([\bar{\mathbf{x}}]_k - z_k) h'(z_k) \mathbf{a}_{col,k}. \quad (4)$$

The proof can be found in Appendix A. Proposition 1 implies that the perturbation introduced by the k th experiment is *limited* only to the *direct out-neighbor* of the perturbed node k . This matches the result in [4], which showed that the influence of a perturbation on a node decays exponentially fast with respect to the shortest distance to the perturbed node. In light of Proposition 1, we let $\delta > 0$ be a pre-defined threshold and consider the index set:

$$\mathcal{S} = \bigcup_{j=1}^K \left\{ (i, j) \in [n] \times [n] : \frac{[\bar{\mathbf{x}} - \mathbf{x}[j]]_i}{[\bar{\mathbf{x}}]_j - z_j} < \delta \right\}. \quad (5)$$

We also define \mathcal{S}_i as the restriction of \mathcal{S} to the i th row of \mathbf{A} , notice that $\mathcal{S}_i \subseteq [K]$. Importantly, \mathcal{S} can be treated as an estimate for the locations of *zeros/non-edges* in the GRN \mathbf{A} .

Let $(\bar{\mathbf{x}}, \{\mathbf{x}[k]\}_{k=1}^K)$ be the *response matrix* in which we gather the gene expression data from K perturbation experiments. When \mathbf{b} is known, we can recover \mathbf{A} by solving for each $i \in [n]$:

$$\min_{\hat{\mathbf{a}}_i} \|\hat{\mathbf{a}}_i\|_1 \quad \text{s.t.} \quad [\bar{\mathbf{x}}]_i = \hat{\mathbf{a}}_i^\top \mathbf{h}(\bar{\mathbf{x}}; \mathbf{b}), \quad (\mathbf{x}[k])_i = \hat{\mathbf{a}}_i^\top \mathbf{h}(\mathbf{x}[k]; \mathbf{b}), \quad \forall k \in [K] \setminus \{i\} \quad (6a)$$

$$[\hat{\mathbf{a}}_i]_i = 0, \quad [\hat{\mathbf{a}}_i]_j = 0, \quad \forall j \in \mathcal{S}_i, \quad \hat{\mathbf{a}}_i \in \mathbb{R}^n. \quad (6b)$$

The solution to the above problem, $\hat{\mathbf{a}}_i$, serves as an estimate for the i th row of \mathbf{A} , i.e., \mathbf{a}_i . We emphasize that the sparsity pattern of the optimum $\hat{\mathbf{a}}_i$, $\{j \in [n] : [\hat{\mathbf{a}}_i]_j \neq 0\}$, can be interpreted as the set of regulators/deregulators to gene i and the ranked list of non-zero elements of $\hat{\mathbf{a}}_i$ indicates the probability of an existing edge in the GRN. The problem can be solved efficiently in parallel using general software package such as `cvx` (Available at <http://www.cvxr.com/cvx>).

2.1 A sufficient condition for *perfect* GRN Recovery

To understand the fundamental limits of recovering the GRN with (6), we study the scenario when the steady states $\bar{\mathbf{x}}$ and $\mathbf{x}[k]$ are obtained *with no noise*, i.e., they satisfy the equalities (2) and (3). The challenge in the analysis is that the *undetermined* linear system (6a) depends on the true network \mathbf{A} itself which is a sparse matrix, and the dynamical system is non-linear. We develop a new sparse recoverability condition given that \mathbf{b} is known. To describe the result, let us state the simplifying assumptions below.

H1 (a) Set \mathcal{S} is the complement of the support of \mathbf{A} ; (b) matrix \mathbf{A} is non-negative; (c) the approximation in Proposition 1 is exact; (d) $\mathbf{h}(\mathbf{x}; \mathbf{b})$ admits an exact first order Taylor approximation; (e) $\bar{x}_k - z_k \geq 0$ for all $k \in [K]$.

Theorem 1 Assume H1. For each $i \in [n]$, if the induced bipartite graph $G(\mathcal{S}_i, [n])$ is an (α, δ) -unbalanced expander graph with left degree bounded in $[d_l, d_u]$ such that

$$(1 + (d_l/d_u)\delta)\|\mathbf{a}_i\|_0 \leq \alpha n, \quad 2(d_l/d_u)\delta > \sqrt{5} - 1, \quad (7)$$

then solving (6) with the additional constraint $\hat{\mathbf{a}}_i \geq 0$ yields a unique solution that $\mathbf{a}_i = \hat{\mathbf{a}}_i$.

The formal definition of an expander bipartite graph can be found in [14, 29] or in Definition 1 of the appendix. The proof of Theorem 1 is relegated to Appendix B. A curious fact about Theorem 1 is that the condition (7) depends on the graph structure of \mathbf{A} as well as the sparsity of each row \mathbf{a}_i of \mathbf{A} . This is due to the fact that the linear system (6a) depend on \mathbf{A} itself.

A possible scenario satisfying the conditions in Theorem 1 requires choosing the set of perturbed genes $[K]$ such that each gene in V is regulated by a similar number ($\sim d$) of the genes in $[K]$. As we shall argue in Appendix B.1, fix $\alpha \in (0, 1)$, when $|\mathcal{S}_i| \geq c(d) \cdot \alpha n$ for some $c(d) > 1$ then $G(\mathcal{S}_i, V)$ is an $(\alpha, 1 - 1/(d - 1))$ -expander with $d_l = d - 1$, $d_u = d$ in high probability. Applying Theorem 1 shows that perfect recovery is possible if

$$K \approx |\mathcal{S}_i| > c(d) \cdot (1 + \delta(d - 1)/d)\|\mathbf{a}_i\|_0 \geq c(d) \cdot ((2d - 2)/d)\|\mathbf{a}_i\|_0, \quad \forall i \in [n], \quad (8)$$

where $\|\mathbf{a}_i\|_0$ counts the number of non-zeros in the vector \mathbf{a}_i . For example, when $d = 10$, $\|\mathbf{a}_i\|_0 = 0.1n$, we have $c(d) \cdot ((2d - 2)/d) \approx 2.91$. Since $\|\mathbf{a}_i\|_0 \leq d_{max}$, where d_{max} is the maximum in-degree, it shows that $K = \Omega(d_{max})$ is a sufficient condition for perfect recovery.

Finally, we remark that the results in Theorem 1 gives a condition for the perfect recovery or identification of \mathbf{A} . This is significantly more challenging than merely *reconstructing* the GRN as required by DREAM5 [21], which focuses only on detecting the edges E of the GRN.

2.2 Robust IDentification of Sparse networks (RIDS) method

So far, the theoretical model above assumes that the steady state expression data are measured noiselessly and the parameter \mathbf{b} in the model response function $h(x; \mathbf{b})$ is known. In practice, \mathbf{b} is unknown and maybe different for every gene. We have the *noisy* expression data:

$$\tilde{\mathbf{x}} = \bar{\mathbf{x}} + \bar{\boldsymbol{\epsilon}} \quad \text{and} \quad \tilde{\mathbf{x}}[k] = \mathbf{x}[k] + \boldsymbol{\epsilon}[k], \quad (9)$$

where the vectors $\bar{\boldsymbol{\epsilon}}, \boldsymbol{\epsilon}[k]$ represent some additive noise of unknown distribution. Let $i \in [n]$, \mathbf{b}_i be the parameter of gene i and define the matrix/vector:

$$\mathbf{y}_i := \begin{pmatrix} \tilde{x}_i[1] \\ \vdots \\ \tilde{x}_i[K] \\ \tilde{\bar{x}}_i \end{pmatrix} \quad \text{and} \quad \mathbf{H}_i(\mathbf{b}_i) := \begin{pmatrix} \mathbf{h}(\tilde{\mathbf{x}}[1]; \mathbf{b}_i)^\top \\ \vdots \\ \mathbf{h}(\tilde{\mathbf{x}}[K]; \mathbf{b}_i)^\top \\ \mathbf{h}(\tilde{\bar{\mathbf{x}}}; \mathbf{b}_i)^\top \end{pmatrix}. \quad (10)$$

Let $\hat{\mathbf{b}}_i$ be an estimate of \mathbf{b}_i . Naturally, one would like to relax the equalities in (6a) and minimize $\lambda\|\hat{\mathbf{a}}_i\|_1 + \|\mathbf{y}_i - \mathbf{H}_i(\hat{\mathbf{b}}_i)\hat{\mathbf{a}}_i\|_2$. However, we observe that the k th element of \mathbf{y}_i is expressed as:

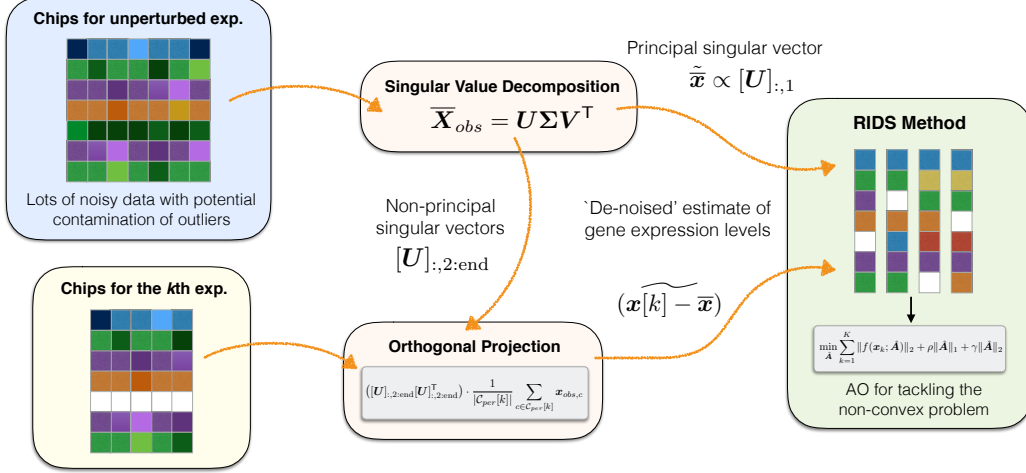


Figure 2: Overview of the proposed RIDS method applied on empirical data. In the pre-processing stage, we first recover the unperturbed steady-state expression levels $\tilde{\mathbf{x}}$ by analyzing the principal component of the stacked response matrix $\bar{\mathbf{X}}_{obs}$. Then, an orthogonal projection is applied to the perturbed steady-state expression levels to recover $\mathbf{x}[k] - \bar{\mathbf{x}}$ for each perturbation condition. This forms $(K + 1)$ vectors where each of them correspond to a distinct perturbation condition (including no perturbation). Finally, we tackle the robust sparse network recovery problem (12) for GRN recovery via the AO procedure (13).

$$\tilde{x}_i[k] = \mathbf{h}(\mathbf{x}[k]; \mathbf{b}_i)^\top \mathbf{a}_i + \epsilon_i[k] = \mathbf{h}(\tilde{\mathbf{x}}[k]; \hat{\mathbf{b}}_i)^\top \mathbf{a}_i + \delta[k]^\top \mathbf{a}_i + \epsilon_i[k], \quad (11)$$

where $\delta[k] := \mathbf{h}(\mathbf{x}[k]; \mathbf{b}_i) - \mathbf{h}(\tilde{\mathbf{x}}[k]; \hat{\mathbf{b}}_i)$ is an unknown vector that scales with the magnitude of $\epsilon[k]$. The difference vector $\mathbf{y}_i - \mathbf{H}_i(\mathbf{b}_i)\hat{\mathbf{a}}_i$ is dependent on \mathbf{a}_i and can not be modeled as an additive noise.

On the other hand, the unknown parameter \mathbf{b}_i lies in a parameter set \mathcal{B} . As such, our robust identification of sparse networks (RIDS) method tackles — for each $i \in [n]$:

$$\min_{\hat{\mathbf{a}}_i, \hat{\mathbf{b}}_i} J(\hat{\mathbf{a}}_i; \hat{\mathbf{b}}_i) := \|\mathbf{y}_i - \mathbf{H}_i(\hat{\mathbf{b}}_i)\hat{\mathbf{a}}_i\|_2 + \rho\|\hat{\mathbf{a}}_i\|_1 + \gamma\|\hat{\mathbf{a}}_i\|_2 \quad (12a)$$

$$\text{s.t. } [\hat{\mathbf{a}}_i]_i = 0, [\hat{\mathbf{a}}_i]_j = 0, \forall j \in \mathcal{S}_i, \hat{\mathbf{a}}_i \in \mathbb{R}^n, \hat{\mathbf{b}}_i \in \mathcal{B}, \quad (12b)$$

where $\rho, \gamma > 0$ are fixed regularization parameter. The derivation details of (12) can be found in Appendix C. This formulation is akin to the matrix uncertainty (MU) selector in [25] for sparse recovery with uncertainty. Despite being robust to measurement error, the above formulation simultaneously solves for the best model parameter that fits with the expression data.

However, (12) is a *non-convex* problem due to the multiplicative coupling in the least square objective function. The problem cannot be solved directly using off-the-shelf packages. The RIDS method applies an alternating optimization (AO) approach to get around with the issue, *i.e.*, by running the following procedure for each $i \in [n]$ — let L be the maximum number of iterations:

$$\begin{aligned} &\text{for } \ell = 1, 2, 3, \dots, L \\ &\quad \hat{\mathbf{a}}_i^{\ell+1} \leftarrow \arg \min_{\hat{\mathbf{a}}_i} J(\hat{\mathbf{a}}_i; \hat{\mathbf{b}}_i^\ell) \text{ s.t. (12b) satisfied,} \\ &\quad \hat{\mathbf{b}}_i^{\ell+1} \leftarrow \arg \min_{\mathbf{b}} \|(\hat{\mathbf{b}}_i^\ell - \epsilon \cdot \nabla_{\mathbf{b}} J(\hat{\mathbf{a}}_i^{\ell+1}; \hat{\mathbf{b}}_i^\ell)) - \mathbf{b}\|_2 \text{ s.t. } \mathbf{b} \in \mathcal{B}, \end{aligned} \quad (13)$$

where $\epsilon > 0$ is a fixed step size and $\nabla_{\mathbf{b}} J(\hat{\mathbf{a}}_i^{\ell+1}; \hat{\mathbf{b}}_i^\ell)$ is the gradient of the cost function. The last step

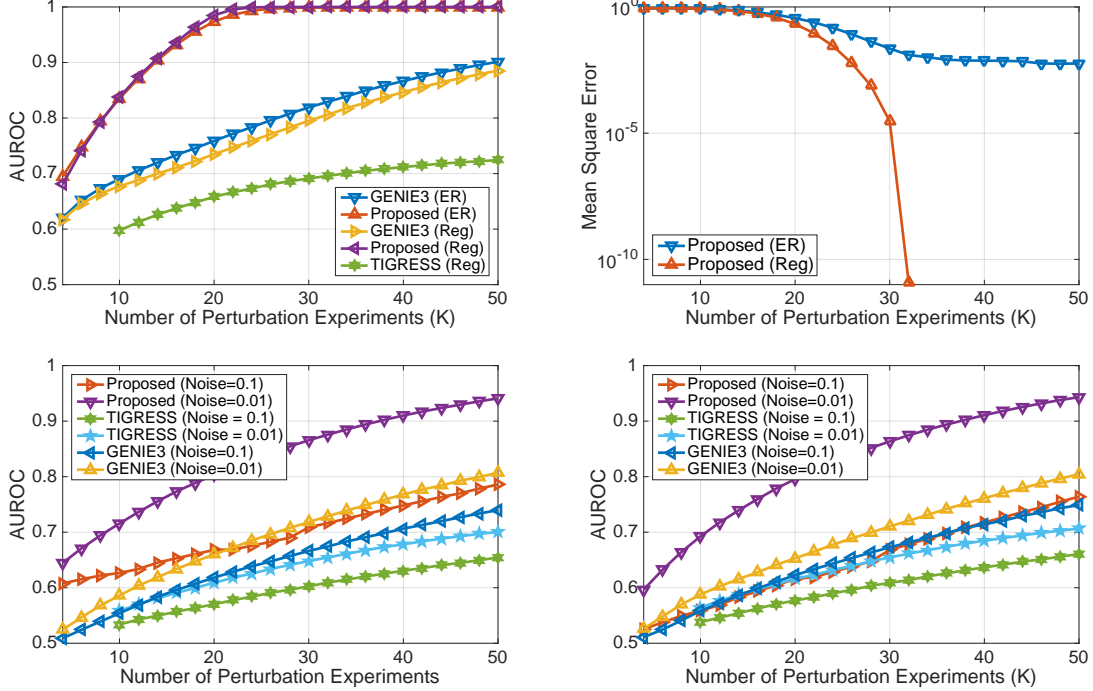


Figure 3: Numerical Experiments on synthetic data of networks with $n = 100$ nodes, averaged over 100 Monte-Carlo trials. Two random graph models are considered — Erdos-Renyi graph with connectivity $p = 0.1$ and Random regular graph with constant degree $d = 10$. The model response function used is $h(x) = x^{0.5}/(1 + x^{0.5})$ and the parameters are assumed to be known. The top figures consider the noiseless observation case. (Top-Left) Area under ROC (AUROC) against the number of perturbation experiments K . (Top-Right) Mean square error of the recovered $\hat{\mathbf{A}}$ using the proposed method. Our proposed RIDS method (via solving (6)) outperforms GENIE3 and TIGRESS over all range of K and achieves perfect recovery at $K \geq 32$ for the random regular graph model (as indicated by the MSE plot). Notice that the TIGRESS method has encountered numerical issues for the case with ER graphs. Perfect recovery was also observed for the ER model for $\sim 70\%$ of the instances at $K \geq 32$, but however, as the ER graphs tend to have hubs with high degree, its average MSE performance will be affected. The bottom figures consider the noisy observation case. (Bottom-Left) Random Erdos-Renyi graphs with connectivity 0.1. (Bottom-Right) Random Regular graphs with constant degree $d = 10$. Again, we observe that the proposed method has outperformed GENIE3 and TIGRESS for the considered range of K .

is a *projected gradient* update step for $\hat{\mathbf{b}}_i$. The RIDS method is summarized in Figure 2. In the first stage, we apply a pre-process method to de-noise the experimental data (see Appendix D); in the second stage, we tackle the robust GRN recovery problem (12) using the AO procedure in (13).

3 Numerical Results & Discussions

3.1 Synthetic Data

In silico data. We test the models when the GRN $G = (V, E)$ with $n = 100$. The weight matrix \mathbf{A} has entries that are uniformly distributed in $(0, 1]$. We evaluate the steady-state gene expression levels subject to gene deletion using the 4th order Runge-Katta method. We include the GENIE3 method [17] and TIGRESS method [16] implemented by their respective authors for benchmarks. The latter two were the best performing methods in the DREAM5 challenge. All algorithms tested

are implemented on MATLAB 2016a. For the zeros index set \mathcal{S} in (5), we set $\delta = 0.02$.

Analysis. Figure 3 (Top) compares the mean square error of the recovered $\hat{\mathbf{A}}$ and the area under an ROC curve (AUROC) for the recovered network G versus the number of perturbation experiments K . We assume noiseless measurements in this case and solve (6) to recover the network. We observe that the proposed method achieves an AUROC of ≥ 0.9 with $K \geq 14$ perturbation experiments, significantly outperforming the GENIE3 and TIGRESS methods under similar conditions. Moreover, we see that the proposed RIDS method has a better performance when the underlying graph is a regular graph. In particular, with $K = 32$ perturbation experiments we have perfect recovery of both the inferred links and interaction strengths, A_{ij} , for the regular graph model. The above observation coincides with Theorem 1 which predicts with $2.91 \times \max_{i \in [n]} \|\mathbf{a}_i\|_0$ perturbation experiments, one can perfectly recover the GRN with our proposed approach when $n \rightarrow \infty$. Nevertheless, having $K \approx 14$ experiments ($\sim 15\%$ of the total number of genes) is sufficient to yield a good GRN reconstruction performance.

Figure 3 (Bottom) considers the *noisy* measurement scenario. With reference to (9), the elements of $\bar{\epsilon}, \epsilon[k]$ are independently extracted from normal distributions $\mathcal{N}(0, 0.1)$ and $\mathcal{N}(0, 0.01)$. We apply the robust formulation (12) with the regularizing parameters set to $\rho = 10^{-5}, \gamma = 0.5, \gamma = 0.05$ for the case with $\mathcal{N}(0, 0.1)$ noise and $\mathcal{N}(0, 0.01)$ noise, respectively, to recover the network, notice that in this scenario \mathbf{b} is known and problem (12) is solved directly. Comparing the average AUROC performance shows that our method has a consistently better performance than GENIE3 and TIGRESS. However, as the noise power grows, the advantage of applying our method declines.

3.2 Empirical Data

In vivo data. To test our methodology against empirical data, we focus on the reconstruction performance of *Escherichia coli* (*E. coli*) and *Saccharomyces cerevisiae* (*S. cerevisiae*) from gene perturbation experiments, using the highly curated datasets collected for the DREAM5 Network Inference Challenge¹. The *E. coli* (*resp. S. cerevisiae*) dataset describes 805 (*resp. 536*) vectors of expression level of $n = 4511$ (*resp. n = 5950*) anonymized genes under different experimental conditions; the dataset also lists a subset of $TF = 334$ (*resp. TF = 333*) genes that are recognized as known *transcription factors* (TFs), some of whom are decoys, namely wrongly labeled as such.

Our method relies on the expression levels in the steady state only. Therefore, we use only 326 out of 805 (*resp. 238* out of 536) vectors of expression levels in the *E. coli* (*resp. S. cerevisiae*) dataset, *i.e.*, about $\sim 40\%$ of the data. The remaining vectors correspond to transient states, which our method is not designed to treat. Upon grouping and denoising the dataset using the pre-processing method in Appendix D, we are left with $K = 56$ (*resp. K = 7*) vectors of expression levels for *E. coli* (*resp. S. cerevisiae*), each with a distinct perturbation conditions. From these vectors, we must reconstruct \mathbf{A} , capturing the GRN between the transcription factors and the genes, a total of 334×4511 (*resp. 333 \times 5950*) potential links for *E. coli* (*resp. S. cerevisiae*).

We use the following model response function for empirical data:

$$h(x; \mathbf{b}) = 0.75 \cdot x^{b_2} / (1 + b_1 x^{b_2}), \quad \mathcal{B} = \{\mathbf{b} \in \mathbb{R}^2 : b_1, b_2 \geq 0\}. \quad (14)$$

The above encompasses several kinetic interaction models. When $b_1 \rightarrow 0$, the interaction model

¹Available at <https://www.synapse.org/#!Synapse:syn2787209/wiki/70351>.

Methods	<i>E. coli</i>			<i>S. cerevisiae</i>		
	AUROC	AUPR	Score	AUROC	AUPR	Score
TIGRESS [16]	0.595	0.069	4.41	0.517	0.02	1.082
GENIE3 [17]	0.617	0.093	14.79	0.518	0.021	1.387
RankSum	0.65	0.09	24.90	0.528	0.022	6.236
bLARS [27]	N/A	N/A	5.841	N/A	N/A	7.479
<i>RIDS</i> (<i>top 100k</i>)	0.6808 1.69×10^{-64}	0.0504 9.9×10^{-2}	32.39	0.525 3.84×10^{-8}	0.022 2.3×10^{-2}	4.694
<i>RIDS</i> (opt. b , <i>top 100k</i>)	0.6823 3.13×10^{-66}	0.0508 8.5×10^{-2}	33.29	0.525 6.47×10^{-8}	0.021 2.9×10^{-2}	4.161
<i>RIDS</i> (no TF, <i>top 100k</i>)	0.6745 2.78×10^{-57}	0.0540 2.0×10^{-2}	29.12	0.524 3.70×10^{-7}	0.0221 4.8×10^{-2}	4.298
iRafNet [23]	0.641	0.112	29.26	N/A	N/A	N/A
GENIMS [31]	0.705	0.052	48.33	0.533	0.02	8.454
<i>RIDS</i> (<i>top 500k</i>)	0.7573 1.04×10^{-184}	0.0574 2.7×10^{-3}	93.28	0.5734 1.5×10^{-119}	0.0252 2.38×10^{-7}	62.64

*All values/scores are calculated with the top 100k predictions. Exceptions are the iRafNet, GENIMS and RIDS (top 500k) in the last three rows, that are based on the top 200k, all, top 500k predictions, respectively.

Table 1: GRN recovery result on the two *in vivo* dataset. Scores for RankSum, GENIE3 and TIGRESS are taken directly from the DREAM5’s leaderboard. Notice that RankSum is the community integrated prediction. For the RIDS method, the lower scores are the p -value for the AUROC/AUPR metrics. The RIDS method in the 5th row uses the median optimized parameters learnt for the two networks, *i.e.*, $b_1 = 0.047, b_2 = 0.5893$ for *E. coli* and $b_1 = 0.5571, b_2 = 0.3749$ for *S. cerevisiae* and we solve (12) with the fixed parameters for all genes; the 6th row corresponds to tackling (12) *without* using the list of transcription factors. Our method obtains the best prediction score for *E. coli*, and the score for *S. cerevisiae* is comparable to state-of-the-art.

tends to be that of a chemical activation process; otherwise the interaction model has a saturating effect close to a switch-like process; b_2 controls the saturation rate of the gene interaction. We apply the RIDS method to learn simultaneously the GRN and the above model parameters. Further details of the experimental procedure can be found in Appendix D.

Analysis. Table 1 shows the GRN recovery result for the truncated top 100,000 and 500,000 predicted links in the GRN, compared to the Gold Standard. For each method, we evaluate the AUROC, AUPR and the prediction score, defined similarly as in [21] by $\text{Score} := (\log_{10}(p_{\text{AUROC}}) + \log_{10}(p_{\text{AUPR}}))/2$, where p_{AUROC} and p_{AUPR} are the p -values for AUROC/AUPR. The numerical experiment demonstrates that our robust GRN recovery approach is able to infer GRN from *few steady-state data*, while achieving superior performance to the state-of-the-art. In particular, for the *E. coli* network, the RIDS method is the top performer among the compared methods, beating even the community integrated prediction (RankSum) while using $\sim 60\%$ less experimental data. Our method gives good performance even the list of TFs are not provided when tackling (12).

The model parameters learnt using RIDS are plotted in Figure 4. We observe that the parameters learnt are clustered around $b_1 \sim 0.05, b_2 \sim 0.6$ for the *E. coli* network and $b_1 \sim 0.4, b_2 \sim 0.55$ for the *S. cerevisiae* network. Suggesting that the interaction model for the former is close to a chemical activation process, while the latter is close to being a combination. The parameters learnt are also similar across different genes, which seems to indicate that the parameter space may be reduced by using a single set of parameters for all genes. This is corroborated by the AUROC/AUPR performance obtained when the RIDS method is re-applied with the same model parameter for all genes in Table 1. Additional results on the empirical data can be found in Appendix E, where we




Figure 4: Model parameters learnt for different networks. (Left) *E. coli* network. (Right) *S. cerevisiae* network. The top figures show the scatter plot of the parameters. The bottom figures show the histograms of parameters learnt, where the blue patch is the saturation rate parameter b_2 and the red patch is parameter b_1 which controls if the interaction model is more of the chemical activation type or the switching type, cf. (14). Meanwhile, the *S. cerevisiae* shows a higher variance for the b_1, b_2 values learnt for each gene, this is due to the fact that the in vivo data available for this network is more scarce than the one for *E. coli*.

further demonstrate that RIDS can infer directionality from empirical GRNs.

Computational Complexity. The problem (6) and the AO procedure in (13) can be solved by applying generic optimization packages such as *cvx*. It takes about 25 seconds to recover the synthetic networks in Section 3.1 with $n = 100$ genes on a quad-core labtop computer and about 1100 seconds for the $n = 4511$ genes *E. Coli*. network in Section 3.2 on a 60-cores server. Further speed up will be possible if we apply specialized optimization tools.

Conclusions. This work proposes the RIDS method for GRN inference using perturbation experiments. The RIDS method is developed through modelling the gene expression data as the outcome of perturbing a nonlinear dynamic system. To improve robustness, the method first applies a subspace projection method for denoising the gene expression data, then a sparse and robust estimator is applied to recover the network. The model parameters of the dynamic system will also be inferred simultaneously. Our theoretical analysis, conducted under the assumption that the dynamic's parameters are known, shows that it is possible to recover the GRN even when there are only a few sets of perturbation experiment data available. This is in contrary to the common belief that it requires a large number of experiments to apply similar methods. Moreover, our experiments on empirical data shows that the RIDS method compares favorably to the state-of-the-art methods while requiring $\sim 60\%$ less data. The RIDS method paves the way to study the dynamics of gene interactions by having the ability to infer the model parameters, which is important as studied by [24]. For example, our preliminary result suggests that the genes in the *E. coli* and *S. cerevisiae* networks tend to have a different interaction model with its neighboring genes.

References

- [1] R. Albert. Scale-free networks in cell biology. *Journal of Cell Science*, 118(21):4947, 10 2005.
- [2] A.-L. Barabasi and Z. N. Oltvai. Network biology: understanding the cell’s functional organization. *Nat Rev Genet*, 5(2):101–113, 02 2004.
- [3] B. Barzel and A.-L. Barabasi. Network link prediction by global silencing of indirect correlations. *Nat Biotech*, 31(8):720–725, 08 2013.
- [4] B. Barzel and A.-L. Barabasi. Universality in network dynamics. *Nat Physics*, 9:673–681, October 2013.
- [5] B. Barzel and O. Biham. Quantifying the connectivity of a network: The network correlation function method. *Physical Review E*, 80(4):046104–, 10 2009.
- [6] R. Bonneau, D. J. Reiss, P. Shannon, M. Facciotti, L. Hood, N. S. Baliga, and V. Thorsson. The inferelator: an algorithm for learning parsimonious regulatory networks from systems-biology data sets de novo. *Genome Biology*, 7(5):R36, 2006.
- [7] M. Buchanan. *Networks in Cell Biology*. Cambridge University Press, 2010.
- [8] X. Cai, J. A. Bazerque, and G. B. Giannakis. Inference of gene regulatory networks with sparse structural equation models exploiting genetic perturbations. *PLoS Comput Biol*, 9(5):e1003068–, 05 2013.
- [9] E. Candes and T. Tao. Decoding by Linear Programming. *IEEE Trans. Inf. Theory*, 51(12):4203–4215, 2005.
- [10] R. De Smet and K. Marchal. Advantages and limitations of current network inference methods. *Nat Rev Micro*, 8(10):717–729, 10 2010.
- [11] N. Friedman, M. Linial, I. Nachman, and D. Pe’er. Using bayesian networks to analyze expression data. *Journal of computational biology*, 7(3-4):601–620, 2000.
- [12] T. S. Gardner, D. di Bernardo, D. Lorenz, and J. J. Collins. Inferring genetic networks and identifying compound mode of action via expression profiling. *Science*, 301(5629):102, 07 2003.
- [13] M. Ghanbari, J. Lasserre, and M. Vingron. Reconstruction of gene networks using prior knowledge. *BMC Systems Biology*, 9(1):84, 2015.
- [14] A. Gilbert and P. Indyk. Sparse recovery using sparse matrices. *Proceedings of the IEEE*, 98(6):937–947, 2010.
- [15] C. T. Harbison, D. B. Gordon, T. I. Lee, N. J. Rinaldi, K. D. Macisaac, T. W. Danford, N. M. Hannett, J.-B. Tagne, D. B. Reynolds, J. Yoo, E. G. Jennings, J. Zeitlinger, D. K. Pokholok, M. Kellis, P. A. Rolfe, K. T. Takusagawa, E. S. Lander, D. K. Gifford, E. Fraenkel, and R. A. Young. Transcriptional regulatory code of a eukaryotic genome. *Nature*, 431(7004):99–104, 09 2004.
- [16] A.-C. Haury, F. Mordélet, P. Vera-Licona, and J.-P. Vert. Tigress: Trustful inference of gene regulation using stability selection. *BMC Systems Biology*, 6(1):1–17, 2012.
- [17] V. A. Huynh-Thu, A. Irrthum, L. Wehenkel, and P. Geurts. Inferring regulatory networks from expression data using tree-based methods. *PLoS ONE*, 5(9), Sept 2010.
- [18] T. Ideker and R. Sharan. Protein networks in disease. *Genome Research*, 18(4):644–652, 04 2008.
- [19] T. Kang, R. Moore, Y. Li, E. Sontag, and L. Bleris. Discriminating direct and indirect connectivities in biological networks. *Proceedings of the National Academy of Sciences*, 112(41):12893–12898, 10 2015.
- [20] R. Küffner, T. Petri, P. Tavakkolkhah, L. Windhager, and R. Zimmer. Inferring gene regulatory networks by anova. *Bioinformatics*, 28(10):1376–1382, 2012.

- [21] D. Marbach, J. C. Costello, R. Kuffner, N. M. Vega, R. J. Prill, D. M. Camacho, K. R. Allison, M. Kellis, J. J. Collins, and G. Stolovitzky. Wisdom of crowds for robust gene network inference. *Nat Meth*, 9(8):796–804, 08 2012.
- [22] N. Omranian, J. M. O. Eloundou-Mbebi, B. Mueller-Roeber, and Z. Nikoloski. Gene regulatory network inference using fused lasso on multiple data sets. *Scientific Reports*, 6:20533 EP –, 02 2016.
- [23] F. Petralia, P. Wang, J. Yang, and Z. Tu. Integrative random forest for gene regulatory network inference. *Bioinformatics*, 31(12):i197–i205, 06 2015.
- [24] M. Ronen, R. Rosenberg, B. I. Shraiman, and U. Alon. Assigning numbers to the arrows: Parameterizing a gene regulation network by using accurate expression kinetics. *Proceedings of the National Academy of Sciences*, 99(16):10555–10560, 08 2002.
- [25] M. Rosenbaum and A. Tsybakov. Sparse recovery under matrix uncertainty. *Annals of Statistics*, 38(5):2620–2651, 2010.
- [26] A. Shojaie, A. Jauhiainen, M. Kallitsis, and G. Michailidis. Inferring regulatory networks by combining perturbation screens and steady state gene expression profiles. *PLoS ONE*, 9(2):e82393–, 02 2014.
- [27] N. Singh and M. Vidyasagar. blars: An algorithm to infer gene regulatory networks. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 13(2):301–314, March 2016.
- [28] H. M. Tran and S. T. Bukkapatnam. Inferring sparse networks for noisy transient processes. *Scientific Reports*, 6:21963, 2016.
- [29] H.-T. Wai, A. Scaglione, and A. Leshem. Active Sensing of Social Networks. *IEEE Trans. Sig. and Inf. Proc. over Networks*, 2(3):406–419, 2016.
- [30] M. Wang, W. Xu, and A. Tang. A unique ”nonnegative” solution to an underdetermined system: From vectors to matrices. *IEEE Trans. Signal Process.*, 59(3):1007–1016, 2011.
- [31] J. Wu, X. Zhao, Z. Lin, and Z. Shao. Large scale gene regulatory network inference with a multi-level strategy. *Molecular BioSystems*, 12(2):588–597, 2016.
- [32] K. Y. Yip, R. P. Alexander, K.-K. Yan, and M. Gerstein. Improved reconstruction of in silico gene regulatory networks by integrating knockout and perturbation data. *PLoS ONE*, 5(1):e8121–, 01 2010.

A Proof of Proposition 1

Define

$$\bar{\mathbf{x}} - \mathbf{x}[k] = ([\bar{\mathbf{x}}]_k - z_k)\mathbf{e}_k + \boldsymbol{\epsilon}, \quad (15)$$

such that the k th component of $\boldsymbol{\epsilon}$ is zero. Our goal is to find $\boldsymbol{\epsilon}$. We have:

$$\bar{\mathbf{x}} - \mathbf{x}[k] = \mathbf{A}(\mathbf{h}(\bar{\mathbf{x}}) - \mathbf{h}(\mathbf{x}[k])) + \mathbf{e}_k \mathbf{e}_k^\top \mathbf{A} \mathbf{h}(\mathbf{x}[k]) + z_k \mathbf{e}_k. \quad (16)$$

Notice that \mathbf{e}_k is in the null space of $(\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top)$. Left-multiplying $(\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top)$ to the both sides of the equation yields:

$$\begin{aligned} \boldsymbol{\epsilon} &= (\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A}(\mathbf{h}(\bar{\mathbf{x}}) - \mathbf{h}(\mathbf{x}[k])) \\ &\approx (\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \nabla \mathbf{h}(\mathbf{x}[k]) ([\bar{\mathbf{x}}]_k - z_k) \mathbf{e}_k + \boldsymbol{\epsilon} \end{aligned} \quad (17)$$

where we have taken the Taylor's expansion for $\mathbf{h}(\bar{\mathbf{x}})$ centered at $\mathbf{x}[k]$ and the approximation is accurate when the perturbation $\bar{\mathbf{x}} - \mathbf{x}[k]$ is small. This gives

$$\boldsymbol{\epsilon} \approx (\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \nabla \mathbf{h}(\mathbf{x}[k]) \boldsymbol{\epsilon} + ([\bar{\mathbf{x}}]_k - z_k) (\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \nabla \mathbf{h}(\mathbf{x}[k]) \mathbf{e}_k \quad (18)$$

Moreover, we notice that $\nabla \mathbf{h}(\mathbf{x}[k])$ is a diagonal matrix with $h'(z_k)$ on its k th entry, therefore the latter term can be simplified as

$$\begin{aligned} ([\bar{\mathbf{x}}]_k - z_k) (\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \nabla \mathbf{h}(\mathbf{x}[k]) \mathbf{e}_k &= ([\bar{\mathbf{x}}]_k - z_k) h'(z_k) (\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{a}_{col,k} \\ &= ([\bar{\mathbf{x}}]_k - z_k) h'(z_k) \mathbf{a}_{col,k}, \end{aligned} \quad (19)$$

where the last equality is due to the fact that the k th element of $\mathbf{a}_{col,k}$ is zero.

Finally, we observe that

$$\begin{aligned} \boldsymbol{\epsilon} &= (\mathbf{I} - (\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \nabla \mathbf{h}(\mathbf{x}[k]))^{-1} \cdot ([\bar{\mathbf{x}}]_k - z_k) h'(z_k) \mathbf{a}_{col,k} \\ &= (\mathbf{I} + (\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \nabla \mathbf{h}(\mathbf{x}[k]) \\ &\quad + ((\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \nabla \mathbf{h}(\mathbf{x}[k]))^2 + \dots) \cdot ([\bar{\mathbf{x}}]_k - z_k) h'(z_k) \mathbf{a}_{col,k} \\ &\approx ([\bar{\mathbf{x}}]_k - z_k) h'(z_k) \mathbf{a}_{col,k}, \end{aligned} \quad (20)$$

where the second equality is due to the Taylor's expansion. Notice that the series expansion holds when $\lambda_{max}((\mathbf{I} - \mathbf{e}_k \mathbf{e}_k^\top) \mathbf{A} \nabla \mathbf{h}(\mathbf{x}[k])) < 1$.

B Proof of Theorem 1

Proof Outline. Using the assumptions stated in the Theorem, we first reduce the linear system into a simple form that involves an underdetermined system with a *sparse* sensing matrix. The sensing matrix is then found to have the same structure/support as the bipartite graph formed by the edges *from* the perturbed nodes *to* all other nodes. Finally, the desired perfect recovery condition is given as a consequence of the expander graph property of this bipartite graph.

Let us give a formal definition of expander graph:

Definition 1 An (α, δ) -unbalanced expander graph, $G(A, B)$, is an A, B -bipartite graph (bigraph) with $|A| = n, |B| = m$ with left degree bounded in $[d_l, d_u]$, i.e., $d(v_i) \in [d_l, d_u]$ for all $v_i \in A$, such that for any $S \subseteq A$ with $|S| \leq \alpha n$, we have $\delta |E(S, B)| \leq |N(S)|$, where $E(S, B)$ is the set of edges from S to B and $N(S) = \{v_j \in B : \exists v_i \in S \text{ s.t. } (v_j, v_i) \in E\}$ is the neighbor set of S in B .

We are ready to begin our proof. Let \mathbf{a}_i be the i th row of \mathbf{A} and define a set of vectors $\{\mathbf{y}_i\}_{i=1}^n$ such that $[\mathbf{y}_i]_k := \bar{x}_i - x_i[k]$. We observe that \mathbf{y}_i is a collection of the data points that depend on \mathbf{a}_i . For simplicity, let us focus on the case when $i \notin [K]$,

$$[\mathbf{y}_i]_k = \mathbf{a}_i^\top (\mathbf{h}(\bar{\mathbf{x}}) - \mathbf{h}(\mathbf{x}[k])) \approx (\bar{x}_k - z_k) \cdot \mathbf{a}_i^\top \nabla \mathbf{h}(\bar{\mathbf{x}}) (h'(z_k) \mathbf{a}_{col,k} + \mathbf{e}_k), \quad (21)$$

where we have applied Proposition 1 and the first order Taylor approximation on $\mathbf{h}(\bar{\mathbf{x}}) - \mathbf{h}(\mathbf{x}[k])$ to yield the result above. Notice that we have dropped the dependence on \mathbf{b} as the ODE parameter is assumed to be known in this setting. After some manipulations and applying the assumptions stated in the theorem, we can express the equation above as

$$\mathbf{y}_i = \begin{pmatrix} \mathbf{\Lambda} & \mathbf{0}_{K \times (n-K)} \end{pmatrix} \mathbf{a}_i + \underbrace{\begin{pmatrix} (\bar{x}_1 - z_1) h'(z_1) \cdot \nabla \mathbf{h}(\bar{\mathbf{x}}) \mathbf{a}_{col,1}^\top \\ \vdots \\ (\bar{x}_K - z_K) h'(z_K) \cdot \nabla \mathbf{h}(\bar{\mathbf{x}}) \mathbf{a}_{col,K}^\top \end{pmatrix}}_{:= \tilde{\mathbf{E}}_i} \mathbf{a}_i, \quad (22)$$

where $\mathbf{\Lambda}$ is an $K \times K$ diagonal matrix with the k th element being $[\mathbf{\Lambda}]_{kk} = \bar{x}_k - z_k$. The challenge is that as the non-zero elements of $\mathbf{\Lambda}$ has a larger magnitude than the matrix $\tilde{\mathbf{E}}_i$ in the latter matrix-vector product, the overall sensing matrix $(\mathbf{\Lambda} \mathbf{0}) + \tilde{\mathbf{E}}_i$ is dominated by the diagonal matrix component. That is, $(\mathbf{\Lambda} \mathbf{0}) + \tilde{\mathbf{E}}_i \approx (\mathbf{\Lambda} \mathbf{0})$. However, this implies that the elements of \mathbf{a}_i over the coordinates $[n] \setminus [K]$ cannot be recovered from \mathbf{y}_i as the rows of the sensing matrix supported only on $[K]$. It is thus necessary to exploit extra information to recover \mathbf{a}_i .

In light of this, we note that $[\mathbf{a}_i]_j$ is zero for all j in \mathcal{S}_i . The first matrix-vector product is thus an $K - |\mathcal{S}_i|$ -sparse vector which is supported on the set $\mathcal{S}_i^c = [K] - \mathcal{S}_i$. As we are interested in studying a sufficient condition for perfect recovery, we see that (22) implies the following linear system with the rows corresponding to \mathcal{S}_i^c removed,

$$[\mathbf{y}_i]_{\mathcal{S}_i} = [\tilde{\mathbf{E}}_i]_{\mathcal{S}_i, \cdot} \mathbf{a}_i. \quad (23)$$

Compared to the original model (22), we can suppress the dominating diagonal component in $\tilde{\mathbf{E}}_i$ using the support knowledge on \mathbf{a}_i .

The remaining task is to verify if the reduced sensing matrix $[\tilde{\mathbf{E}}_i]_{\mathcal{S}_i, \cdot}$ is a good sensing matrix. Notice that *dense* and random matrices are known to exhibit good properties in which one only requires $|\mathcal{S}_i| \geq 2 \|\mathbf{a}_i\|_0 \cdot \log n$ to achieve perfect recovery. On the other hand, $[\tilde{\mathbf{E}}_i]_{\mathcal{S}_i, \cdot}$ is a *sparse* matrix whose support depends on the out-neighbors of the nodes in \mathcal{S}_i . In particular, we have $\text{supp}([\tilde{\mathbf{E}}_i]_{\mathcal{S}_i, \cdot}) = \text{supp}(\mathbf{A}_{\cdot, \mathcal{S}_i}^\top)$. An interesting observation is that the support of the sensing matrix depends on the support of \mathbf{A} itself or the network that we wish to recover.

As it turns out, the perfect recovery condition for \mathbf{a}_i boils down to studying conditions on the *support* of $[\tilde{\mathbf{E}}_i]_{\mathcal{S}_i, \cdot}$. In particular, let $G_{bi}(A, B)$ with $|A| = n$ and $|B| = |\mathcal{S}_i| \leq K$ be the bi-partite graph representation of $[\mathbf{E}_i]_{\mathcal{S}_i, \cdot}$ such that the adjacency matrix of G_{bi} , i.e., $\mathbf{A}_{bi} \in \mathbb{R}^{|\mathcal{S}_i| \times n}$, has the same support as $[\mathbf{E}_i]_{\mathcal{S}_i, \cdot}$.

Theorem 2 *If G_{bi} is an (α, δ) -unbalanced expander graph with left degree bounded in $[d_l, d_u]$ such that $2(d_l/d_u) \cdot \delta > \sqrt{5} - 1$ and $k \leq \frac{\alpha}{1+\rho\delta}n$, and $\tilde{\mathbf{E}}_i$ or $-\tilde{\mathbf{E}}_i$ is non-negative, then the set*

$$\mathcal{C} = \{\hat{\mathbf{a}}_i : [\tilde{\mathbf{E}}_i]_{\mathcal{S}_i, \cdot}(\hat{\mathbf{a}}_i - \mathbf{a}_i) = \mathbf{0}\} \quad (24)$$

is a singleton for all k -sparse vector \mathbf{a}_i .

The proof of the theorem above can be found in subsection B.2. Consequently, we observe that all the conditions in Theorem 2 hold, therefore solving (6) yields $\hat{\mathbf{A}} = \mathbf{A}$, i.e., we achieve perfect recovery of the network.

B.1 A special case satisfying the conditions in Theorem 1

Let us first borrow the following proposition from [29]:

Proposition 2 [29, Proposition 5] *Let $G(A, B)$ be a random bipartite with $|A| = n$, $|B| = \beta \cdot n$ for some $\beta < 1$ and the left degree bounded in $[d_l, d_u]$ (i.e., nodes in A has bounded degree). Fix $\alpha \in (0, 1)$ and the following holds:*

$$d_l \geq 3, \quad \beta > \alpha, \quad d_l > (H(\alpha) + \beta H(\alpha/\beta))/(\alpha \log(\beta/\alpha)), \quad (25)$$

where $H(\alpha) = \alpha \log \alpha + (1 - \alpha) \log(1 - \alpha)$ is the binary entropy function. Then, with high probability as $n \rightarrow \infty$, the graph $G(A, B)$ is an $(\alpha, 1 - 1/d_l)$ -expander.

Following our discussions in Section 2.1, we suppose that the perturbed genes $[K]$ are chosen such that for each $j \in V$, the j th gene has d in-neighbors in $[K]$. Notice that this requires selecting perturbation nodes strategically, e.g., using advices from the experts.

Furthermore, the bipartite graph of the $K \times n$ matrix $([\mathbf{A}]_{:, [K]})^\top$ may follow a regular random graph model with constant degree d . Then, using similar arguments as in [29, Proposition 4], it can be shown that the support of the sub-matrix $([\mathbf{A}]_{:, \mathcal{S}_i})^\top$ corresponds to a random bipartite graph with bounded degree in $[d - 1, d]$ with high probability. The intuition is that the matrix $([\mathbf{A}]_{:, \mathcal{S}_i})^\top$ is formed by deleting a small number of rows from the $K \times n$ matrix $([\mathbf{A}]_{:, [K]})^\top$.

Now fix $\alpha \in (0, 1)$. Applying Proposition 2 shows that there exists $\beta = c(d) \cdot \alpha$ with $c(d) > 1$ such that this bipartite graph is an $(\alpha, 1 - 1/(d - 1))$ -expander with high probability and β satisfying the conditions in (25). Taking $|\mathcal{S}_i| = \beta n > c(d) \cdot \alpha n$, we can substitute the numbers into the conditions (7) required by Theorem 1 to obtain an upper bound on $\|\mathbf{a}_i\|_0$ where we can recover \mathbf{A} perfectly. This yields the desirable result in (8).

B.2 Proof of Theorem 2

To prove Theorem 2, the following Lemma would be instrumental. Denote $\text{Null}(\mathbf{E})$ as the null space of a matrix \mathbf{E} , $S_-(\mathbf{w}) = \{i \in [n] : w_i < 0\}$ and $S_+(\mathbf{w}) = \{i \in [n] : w_i > 0\}$ be the negative and positive support of \mathbf{w} .

Lemma 1 *If (i) $\bar{\mathbf{x}}$ is k -sparse, (ii) $\mathbf{E} \in \mathbb{R}^{m \times n}$ satisfies that $\mathbf{0} \neq \mathbf{w} \in \text{Null}(\mathbf{E})$, $|S_-(\mathbf{w})| \geq k + 1$, then the set $\mathcal{C} = \{\mathbf{x} : \mathbf{E}(\mathbf{x} - \bar{\mathbf{x}}) = \mathbf{0}, \mathbf{x} \geq \mathbf{0}\}$ is a singleton.*

Proof. Suppose $|\mathcal{C}| > 1$ such that there exists $\tilde{\mathbf{x}} \in \mathcal{C}$, $\tilde{\mathbf{x}} \neq \bar{\mathbf{x}}$. It is straightforward to show that $\tilde{\mathbf{x}} = \bar{\mathbf{x}} + \mathbf{w}$, where $\mathbf{w} \in \text{Null}(\mathbf{E})$. The assumption implies that $|S_-(\mathbf{w})| \geq k+1$ and $\tilde{\mathbf{x}} = \bar{\mathbf{x}} + \mathbf{w} \not\leq \mathbf{0}$ as $\bar{\mathbf{x}}$ is k -sparse. This contradicts $\tilde{\mathbf{x}} \in \mathcal{C}$. **Q.E.D.**

The next step is to apply a generalization of [30, Theorem 4]:

Lemma 2 *Let $n > m$ and $\mathbf{E} \in \mathbb{R}^{m \times n}$ be a non-negative matrix that has the same support as the adjacency matrix of an (α, δ) -unbalanced bipartite expander graph with left degrees bounded in $[d_l, d_u]$ and $\rho = d_l/d_u$. If $\rho\delta > (\sqrt{5} - 1)/2$, then for all k -sparse vector $\bar{\mathbf{x}}$. The set*

$$\mathcal{C} = \{\mathbf{x} : \mathbf{E}(\mathbf{x} - \bar{\mathbf{x}}) = \mathbf{0}, \mathbf{x} \geq \mathbf{0}\}, \quad (26)$$

is a singleton if $k \leq \frac{\alpha}{1+\rho\delta}n$.

Proof. Using Lemma 1, it suffices to prove that any $\mathbf{w} \in \text{Null}(\mathbf{E})$ with $S_-(\mathbf{w}) \subseteq A$, we have $|S_-(\mathbf{w})| \geq k+1$. We shall proceed by contradiction. Suppose that there exists $\mathbf{w} \in \text{Null}(\mathbf{E})$ such that $|S_-(\mathbf{w})| \leq k$. Since $|S_-(\mathbf{w})| \leq k \leq \alpha n$, the expander property implies:

$$\delta d_l \cdot |S_-(\mathbf{w})| \leq \delta |E(S_-(\mathbf{w}), B)| \leq |N(S_-(\mathbf{w}))|, \quad (27)$$

where $N(S)$ denotes the set of neighbors of S and $E(A, B)$ is the set of edges between A and B . The right hand side can be further upper bounded as:

$$|N(S_-(\mathbf{w}))| \leq |E(S_-(\mathbf{w}), B)| \leq d_u \cdot |S_-(\mathbf{w})| \quad (28)$$

Moreover, we know that $N(S_-(\mathbf{w})) = N(S_+(\mathbf{w})) = N(S_-(\mathbf{w}) \cup S_+(\mathbf{w}))$. This is because $\mathbf{E}\mathbf{w} = \mathbf{0}$ and \mathbf{A} is non-negative, thus the neighborhood sets must coincide to enforce nullity, otherwise this will result in $\mathbf{E}\mathbf{w} \neq \mathbf{0}$. Note that in [30], the proof is achieved by assuming that \mathbf{E} is the binary adjacency matrix. We extend the same argument to the case when \mathbf{E} is non-negative. Using the above discussions and the inequality (27) applied on $N(S_+(\mathbf{w})) = N(S_-(\mathbf{w}))$, we have:

$$|S_+(\mathbf{w})| \geq |N(S_+(\mathbf{w}))|/d_u \geq (d_l/d_u)\delta |S_-(\mathbf{w})| = \rho\delta |S_-(\mathbf{w})|. \quad (29)$$

As $S_-(\mathbf{w})$ and $S_+(\mathbf{w})$ are disjoint, we have $|S_+(\mathbf{w}) \cup S_-(\mathbf{w})| \geq (1 + \rho\delta)|S_-(\mathbf{w})|$. As such, we choose an arbitrary subset $\tilde{S} \subseteq S_+(\mathbf{w}) \cup S_-(\mathbf{w})$ such that $|\tilde{S}| = (1 + \rho\delta)|S_-(\mathbf{w})|$. Notice that $|\tilde{S}| = (1 + \rho\delta)|S_-(\mathbf{w})| \leq \alpha n$. Using the expander property again gives:

$$|N(\tilde{S})| \geq \delta |E(\tilde{S}, B)| \geq d_l \delta (1 + \rho\delta) |S_-(\mathbf{w})| > d_u |S_-(\mathbf{w})| \quad (30)$$

As $\rho\delta > (\sqrt{5} - 1)/2$, the last inequality is valid as $\rho\delta(1 + \rho\delta) > 1$.

Finally, we reach a contradiction as

$$|N(\tilde{S})| \leq |N(S_-(\mathbf{w}) \cup S_+(\mathbf{w}))| = |N(S_-(\mathbf{w}))| \leq d_u |S_-(\mathbf{w})|, \quad (31)$$

leading to $d_u |S_-(\mathbf{w})| > d_u |S_-(\mathbf{w})|$. The lemma is thus proven. **Q.E.D.**

Identifying that G_{bi} satisfies the conditions in Lemma 2, our desirable results can be obtained.

C Derivation of Problem (12)

From (11), for each $i \in [n]$, we can model the vector \mathbf{y}_i as

$$\mathbf{y}_i = \mathbf{H}_i \mathbf{a}_i + \mathbf{\Delta}_i \mathbf{a}_i + \boldsymbol{\epsilon}, \quad (32)$$

where $\mathbf{\Delta}_i$ models the error in the observed measurement matrix \mathbf{H}_i . Let $r > 0$, the uncertainty set for $\mathbf{\Delta}_i$ is defined such that each row vector in the matrix has a bounded norm, *i.e.*,

$$\mathcal{U}_r = \{\mathbf{\Delta}_i : \|\mathbf{d}_k\|_2 \leq r, \forall k \in [K+1]\}, \quad (33)$$

where \mathbf{d}_k is the k th row vector of $\mathbf{\Delta}_i$.

To recover \mathbf{a}_i , we consider minimizing the following robust objective:

$$J(\hat{\mathbf{a}}_i) = \lambda \|\hat{\mathbf{a}}_i\|_1 + \max_{\mathbf{\Delta}_i \in \mathcal{U}_r} \|\mathbf{y}_i - \mathbf{H}_i \hat{\mathbf{a}}_i - \mathbf{\Delta}_i \hat{\mathbf{a}}_i\|_2, \quad (34)$$

which can be upper bounded by:

$$\begin{aligned} J(\hat{\mathbf{a}}_i) &\leq \lambda \|\hat{\mathbf{a}}_i\|_1 + \|\mathbf{y}_i - \mathbf{H}_i \hat{\mathbf{a}}_i\|_2 + \max_{\mathbf{\Delta}_i \in \mathcal{U}_r} \|\mathbf{\Delta}_i \hat{\mathbf{a}}_i\|_2 \\ &= \lambda \|\hat{\mathbf{a}}_i\|_1 + \|\mathbf{y}_i - \mathbf{H}_i \hat{\mathbf{a}}_i\|_2 + r\sqrt{K+1} \cdot \|\hat{\mathbf{a}}_i\|_2, \end{aligned} \quad (35)$$

where the last equality is achieved by applying Cauchy-Schwarz and setting each row of $\mathbf{\Delta}_i$ to $r \cdot \hat{\mathbf{a}}_i / \|\hat{\mathbf{a}}_i\|_2$. Setting $\gamma = r\sqrt{K+1}$ and minimizing the upper bound function yields the robust network recovery problem (12).

D Details of experiments on empirical data

This section describes the pre/post-processing procedures for the empirical support data.

Subspace projection for ‘denoising’. We apply a subspace projection method as a pre-processing stage in the RIDS method. In particular, let the set of chips (or vectors of expression levels) taken under the *no perturbation* condition be \mathcal{C}_{nopt} , we can write the gene expression levels from the c th chip as

$$\bar{\mathbf{x}}_{obs,c} = \bar{\mathbf{x}} + \mathbf{w}_c, \forall c \in \mathcal{C}_{nopt} \implies \underbrace{(\dots \bar{\mathbf{x}}_{obs,c} \dots)}_{:=\bar{\mathbf{X}}_{obs}} = \bar{\mathbf{x}} \mathbf{1}^\top + \mathbf{W}_{obs}, \quad (36)$$

where $\bar{\mathbf{x}}$ is the unperturbed steady state satisfying $\bar{\mathbf{x}} = \mathbf{A}\mathbf{h}(\bar{\mathbf{x}})$ (cf. Eq. (2)) and \mathbf{w}_c is modeled as an additive noise. When the noise is small, we observe that the $n \times |\mathcal{C}_{nopt}|$ matrix $\bar{\mathbf{X}}_{obs}$ formed by stacking up $\bar{\mathbf{x}}_{obs,c}$ horizontally is close to rank-one. In light of this, we recover $\bar{\mathbf{x}}$ by taking the top left-singular vector of $\bar{\mathbf{X}}_{obs}$, *i.e.*,

$$\hat{\bar{\mathbf{x}}} = \sigma_1(\bar{\mathbf{X}}_{obs}) \cdot \mathbf{u}_1 \quad \text{where} \quad \bar{\mathbf{X}}_{obs} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top, \quad (37)$$

and $\sigma_1(\bar{\mathbf{X}}_{obs})$ is the largest singular value of $\bar{\mathbf{X}}_{obs}$. Notice that $\hat{\bar{\mathbf{x}}}$ lies in the null space of the sub-matrix $[\mathbf{U}]_{:,2:\text{end}}$. For each $k \in [K]$, we can estimate the difference $\mathbf{x}[k] - \bar{\mathbf{x}}$ by

$$\widetilde{(\mathbf{x}[k] - \bar{\mathbf{x}})} = ([\mathbf{U}]_{:,2:\text{end}}[\mathbf{U}]_{:,2:\text{end}}^\top)^{-1} \cdot \frac{1}{|\mathcal{C}_{per}[k]|} \sum_{c \in \mathcal{C}_{per}[k]} \mathbf{x}_{obs,c}, \quad (38)$$

where $\mathbf{x}_{obs,c}$ is the gene expression levels from chip c and $\mathcal{C}_{per}[k]$ is the set of chips that corresponds to the k th perturbation experiment. Naturally, we have $\tilde{\mathbf{x}}[k] = (\mathbf{x}[k] - \bar{\mathbf{x}}) + \tilde{\mathbf{x}}$.

Normalizing the gene expression data. The vectors of gene expression level $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{x}}[k]$ obtained from the above are normalized through dividing the vectors by the constant

$$c_{norm} := \frac{1}{2} \max \left\{ \max_{i \in [n]} (\tilde{x}_i), \max_{k \in [K]} \max_{i \in [n]} (\tilde{x}_i[k]) \right\} \quad (39)$$

such that the values of the normalized gene expression level ranges in $[0, 2]$.

Parameters of RIDS. For the zeros index set \mathcal{S} in (5), we set $\delta = 0.005$. The regularization parameters ρ and γ are set to 3×10^0 and 5×10^0 (*resp.* 3×10^{-1} and 5×10^{-1}) in (12) for *E. coli* (*resp.* *S. cerevisiae*) network. Meanwhile, the AO procedure (13) aims at tackling (12) with the model response function template (14), which has several parameters to be initialized — for each $i \in [n]$, we initialize the AO procedure by setting $b_1 = 0.5$ and $b_2 = 0.5$ and the step size ϵ is set to 1×10^{-2} . Lastly, the AO procedure terminates after $L = 10$ iterations in the numerical experiment for a better trade-off between complexity and accuracy.

Post-processing of the estimated GRN. The proposed algorithm may recover also the negative edge weights in $\hat{\mathbf{A}}$. However, we shall treat these negative edge weights as positive ones for the ease of benchmarking using the DREAM5 evaluation script. In particular, we consider $|\hat{\mathbf{A}}|$ as our estimate of the GRN. Furthermore, we normalize the elements in $|\hat{\mathbf{A}}|$ by $\max_{i,j} |[\hat{\mathbf{A}}]_{ij}|$ such that they range in $[0, 1]$. For the AUROC/AUPR evaluation, we only count the top 100,000 (and 500,000) predicted links made by our algorithm.

E Additional Results on Empirical Data

We present additional experiment results for the *in vivo* dataset from DREAM5.

E.1 GRN recovery without the aid of transcription factors

For simplicity, we shall only focus on the *E. coli* network. We consider applying the RIDS method without the help of the information provided from the list of transcription factors. In particular, we solve (12) by fixing the model parameter at $\mathbf{b} = (0.047, 0.5893)$ and set $\rho = 3$, $\gamma = 5$.

AUROC and AUPR scores. Even without any aid from the list of transcription factors, the RIDS method continues to yield a top AUROC/AUPR performance. In particular, the top 100,000 predicted links from the identified $\hat{\mathbf{A}}$, restricted to the DREAM5 identified transcription factors, yields an AUROC of 0.6745 and AUPR of 0.0540 for the *E. coli* network, as seen in Table 1.

Directionality of the GRN. We test whether the directionality in GRNs can be inferred under the scenario described in this subsection *i.e.*, without using the provided list of transcription factors (TFs). To emphasize on the links with higher weights, we normalize the inferred GRN $\hat{\mathbf{A}}$ as $[\tilde{\mathbf{A}}]_{ij} := |\hat{A}_{ij}|^q / \sum_{i',j'} |\hat{A}_{i'j'}|^q$ for $q = 2$. Then we compare the following for each gene $i \in [n]$ —

- Regulatory strengths — absolute sum of the columns of $\tilde{\mathbf{A}}$, *i.e.*, $\sum_{j=1}^n |\tilde{A}_{ji}|$.
- Strengths to be regulated — absolute sum of the rows of $\tilde{\mathbf{A}}$, *i.e.*, $\sum_{j=1}^n |\tilde{A}_{ij}|$.




Figure 5: ROC and Precision-recall curves for the *in vivo* networks. (Top) *E. coli* network. (Bottom) *S. cerevisiae* network. These curves are produced using the DREAM5 evaluation script and the AUROC/AUPR values can be found in Table 1.

The comparison can be found in Figure 6. We observe that the inferred GRN is directional as fewer genes have significantly higher regulatory strengths exerted on the others, while the regulatory strengths received for the genes are more uniformly distributed.




Figure 6: Inferring directionality of the GRN using RIDS — comparison of the inferred regulatory strengths and the received regulatory strengths for each gene in *E. coli*. The top figure shows the regulatory strengths for each gene exerted on the other genes; the bottom figure shows the regulatory strengths *received* for each gene from the other genes.